arXiv:1703.05537v1 [cs.LG] 16 Mar 2017

Shift Aggregate Extract Networks

Francesco Orsini [1 2] **Daniele Baracchi** [2] **Paolo Frasconi** [2]

Abstract

We introduce an architecture based on deep hierarchical decompositions to learn effective representations of large graphs. Our framework extends classic $\mathcal{R}$-decompositions used in kernel methods, enabling nested *part-of-part* relations. Unlike recursive neural networks, which unroll a template on input graphs directly, we unroll a neural network template over the decomposition hierarchy, allowing us to deal with the high degree variability that typically characterize social network graphs. Deep hierarchical decompositions are also amenable to domain compression, a technique that reduces both space and time complexity by exploiting symmetries. We show empirically that our approach is competitive with current state-of-the-art graph classification methods, particularly when dealing with social network datasets.

1. Introduction

Structured data representations are common in application domains such as chemistry, biology, natural language, and social network analysis. In these domains, one can formulate a supervised learning problem where the input portion of the data is a graph, possibly with attributes on vertices and edges. While learning with graphs of moderate size (tens up to a few hundreds of nodes) can be afforded with many existing techniques, scaling up to large networks poses new significant challenges that still leave room for improvement, both in terms of predictive accuracy and in terms of computational efficiency.

Devising suitable representations for graph learning is crucial and nontrivial. A large body of literature exists on the subject, where graph kernels (GKs) and recurrent neural networks (RNNs) are among the most common approaches. GKs follow the classic R-decomposition approach of Haussler (1999). Different kinds of substructures (e.g., shortest-paths (Borgwardt & Kriegel, 2005), graphlets (Shervashidze et al., 2009) or neighborhood subgraph pairs (Costa & De Grave, 2010)) can used to compute the similarity between two graphs in terms of the similarities of the respective sets of parts. RNNs (Sperduti & Starita, 1997; Goller & Kuchler, 1996; Scarselli et al., 2009) unfold a template (with shared weights) over each input graph and construct the vector representation of a node by recursively composing the representations of its neighbors. These representations are typically derived from a loss minimization procedure, where gradients are computed by the backpropagation through structure algorithm (Goller & Kuchler, 1996). One advantage of RNNs over GKs is that the vector representations of the input graphs are learned rather than handcrafted.

Most of GK- and RNN-based approaches have been applied to relatively small graphs, such as those derived from molecules (Ralaivola et al., 2005; Bianucci et al., 2000; Borgwardt & Kriegel, 2005), natural language sentences (Socher et al., 2011) or protein structures (Vullo & Frasconi, 2004; Baldi & Pollastri, 2003; Borgwardt et al., 2005b). On the other hand, large graphs (especially social networks) typically exhibit a highly skewed degree distribution that originates a huge vocabulary of distinct subgraphs. This scenario makes finding a suitable representation much harder: kernels based on subgraph matching would suffer diagonal dominance (Schoelkopf et al., 2002), while RNNs would face the problem of composing a highly variable number of substructure representations in the recursive step. A recent work by Yanardag & Vishwanathan (2015) proposes deep graph kernels (DGK) to upgrade existing graph kernels with a feature reweighing schema that employs CBOW/Skip-gram embedding of the substructures. Another recent work by Niepert et al. (2016) casts graphs into a format suitable for learning with convolutional neural networks (CNNs). These methods have been applied successfully to small graphs but also to graphs derived from social networks.

In this paper, we introduce a novel architecture for supervised graph learning, called shift-aggregate-extract network (SAEN). The architecture operates on hierarchical decompositions of structured data. Like the flat $\mathcal{R}$-decompositions commonly used to define kernels on structured data (Haussler, 1999), $\mathcal{H}$-decompositions are based

[1]University of Leuven, Belgium [2]Università degli Studi di Firenze, Italy. Correspondence to: Francesco Orsini <francesco.orsini@kuleuven.be>.

.

on the *part-of* relation, but allow us to introduce a deep recursive notion of *parts of parts*. At the top level of the hierarchy lies the *whole* data structure. Objects at each intermediate level are decomposed into parts that form the subsequent level of the hierarchy. The bottom level consists of atomic objects, such as, for example, individual vertices or edges of a graph.

SAEN compensates some limitations of recursive neural networks by adding two synergetic degrees of flexibility. First, it unfolds a neural network over a hierarchy of parts rather than using the edge set of the input graph directly; this makes it easier to deal with very high degree vertices. Second, it imposes weight sharing and fixed size of the learned vector representations on a per level basis instead of globally; in this way, more complex parts may be embedded into higher dimensional vectors, without forcing to use excessively large representations for simpler parts.

A second contribution of this paper is a *domain compression* algorithm that can significantly reduce memory usage and runtime. It leverages mathematical results from lifted linear programming (Mladenov et al., 2012) in order to exploit symmetries and perform a lossless compression of $\mathcal{H}$-decompositions.

The paper is organized as follows. In Section 2 we introduce $\mathcal{H}$-decompositions, a generalization of Haussler's (Haussler, 1999) $\mathcal{R}$-decomposition relations. In Section 3 we describe SAEN, a neural network architecture for learning vector representations of $\mathcal{H}$-decompositions. Furthermore, in Section 4 we explain how to exploit symmetries in $\mathcal{H}$-decompositions in order to reduce memory usage and runtime. In Section 5 we report experimental results on several number of real-world datasets. Finally, in Section 6 we discuss some related works and draw some conclusions in Section 7.

2. $\mathcal{H}$-decompositions

We define here a deep hierarchical extension of Haussler's (Haussler, 1999) $\mathcal{R}$-decomposition relation. An $\mathcal{H}$-decomposition represents structured data as a hierarchy of π-parametrized parts. It is formally defined as the triple $(\{S_l\}_{l=0}^{L}, \{\mathcal{R}_{l,\pi}\}_{l=1}^{L}, X)$ where:

- $\{S_l\}_{l=0}^{L}$ are disjoint sets of objects S_l called levels of the hierarchy. The bottom level S_0 contains atomic (i.e. non-decomposable) objects, while the other levels $\{S_l\}_{l=1}^{L}$ contain compound objects, $s \in S_l$, whose parts $s' \in S_{l-1}$ belong to the preceding level, S_{l-1}.
- $\{\mathcal{R}_{l,\pi}\}_{l=1}^{L}$ is a set of l,π-parametrized $\mathcal{R}_{l,\pi}$-convolution relations. Where $\pi \in \Pi_l$ is a membership type from a finite alphabet Π_l of size $n(l) = |\Pi_l|$. At the bottom level, $n(0) = 1$. A pair $(s, s') \in S_l \times S_{l-1}$ belongs to $\mathcal{R}_{l,\pi}$ iff s' is part of s with membership type π. For notational convenience, the parts of s are denoted as $\mathcal{R}_{l,\pi}^{-1}(s) = \{s' | (s', s) \in \mathcal{R}_{l,\pi}\}$.
- X is a set $\{\mathbf{x}(s)\}_{s \in S_0}$ of p-dimensional vectors of attributes assigned to the elements s the bottom layer S_0.

The membership type π is used to represent the roles of the parts of an object. For example, we could decompose a graph as a multiset of π-neighborhood subgraphs [1] in which π is the radius of the neighborhoods. Another possible use of the π-membership type is to distinguish the root from the other vertices in a rooted neighborhood subgraph. Both uses π-membership type are shown in Figure 2.

An $\mathcal{H}$-decomposition is a multilevel generalization of $\mathcal{R}$-convolution relations, and it reduces to an $\mathcal{R}$-convolution relation for $L = 1$.

For example we could produce a 4-levels decomposition by decomposing graph $G \in S_3$ into a set of r-neighborhood subgraphs $g \in S_2$ and employ their radius r as membership type. Furthermore, we could extract shortest paths $p \in S_1$ from the r-neighborhood subgraphs and use their length as membership type. Finally, each shortest path could be decomposed in vertices $v \in S_0$ using their index in the shortest path as membership type.

Another example of decomposition could come from text processing, documents $d \in S_3$ could be decomposed in sentences $s \in S_2$ which are themselves represented as graphs of dependency relations and further decomposed as bags of shortest paths $p \in S_1$ in the dependency graph. Finally, the words $w \in S_0$ (which are the vertices of the dependency graph) constitute the bottom layer and can be represented in attributed form as word vectors.

3. Learning representations with SAEN

A shift-aggregate-extract network (SAEN) is a composite function that maps objects at level l of an $\mathcal{H}$-decomposition into $d(l)$-dimensional real vectors. It uses a sequence of parametrized functions $\{f_0, \ldots, f_L\}$, for example a sequence of neural networks with parameters $\theta_0, \ldots, \theta_L$. At each level, $l = 0, \ldots, L$, each function $f_l : \mathbb{R}^{n(l)d(l)} \rightarrow \mathbb{R}^{d(l+1)}$ operates as follows:

1. It receives as input the *aggregate* vector $\mathbf{a}_l(s)$ de-

[1] The r-neighborhood subgraph (or ego graph) of a vertex v in a graph G is the induced subgraph of G consisting of all vertices whose shortest-path distance from v is at most r.

fined as:

$$\mathbf{a}_l(s) = \begin{cases} \mathbf{x}(s) & \text{if } l = 0 \\ \sum_{\pi \in \Pi_l} \sum_{s' \in \mathcal{R}_{l,\pi}^{-1}(s)} \mathbf{z}_\pi \otimes \mathbf{h}_{l-1}(s') & \text{if } l > 0 \end{cases} \quad (1)$$

where $\mathbf{x}(s)$ is the vector of attributes for object s.

2. It *extracts* the vector representation of s as

$$\mathbf{h}_l(s) = f_l(\mathbf{a}_l(s); \theta_l). \quad (2)$$

The vector $\mathbf{a}_l(s)$ is obtained in two steps: first, previous level representations $\mathbf{h}_{l-1}(s')$ are *shifted* via the Kronecker product $\otimes$ using an indicator vector $\mathbf{z}_\pi \in \mathbb{R}^{n(l)}$. This takes into account of the membership types π. Second, shifted representations are *aggregated* with a sum. Note that all representation sizes $d(l)$, $l > 0$ are hyperparameters that need to be chosen or adjusted.

The shift and aggregate steps are identical to those used in kernel design when computing the explicit feature of a kernel $k(x, z)$ derived from a sum $\sum_{\pi \in \Pi} k_\pi(x, z)$ of base kernels $k_\pi(x, z)$, $\pi \in \Pi$. In principle, it would be indeed possible to turn SAEN into a kernel method by removing the extraction step. However, such an approach would increase the dimensionality of the feature space by a multiplicative factor $n(l)$ for each level l of the $\mathcal{H}$-decomposition, thus leading to an exponential number of features. When using SAEN, the feature space growth is prevented by exploiting a distributed representation (via a multilayered neural network) during the extraction step. As a result, SAEN can easily cope with $\mathcal{H}$-decompositions consisting of multiple levels.

4. Exploiting symmetries for domain compression

In this section we propose a technique, called *domain compression*, which allows us to save memory and speedup the SAEN computation. Domain compression exploits symmetries in $\mathcal{H}$-decompositions by collapsing equivalent objects in each level. Domain compression requires that the attributes $\mathbf{x}(s)$ of the elements s in the bottom level S_0 are categorical.

Two objects a, b in a level S_l are *collapsible*, denoted $a \sim b$, if they share the same representation, i.e., $\mathbf{h}_l(a) = \mathbf{h}_l(b)$ for all the possible values of $\theta_0, \ldots, \theta_l$. A compressed level S_l^{comp} is the quotient set of level S_l with respect to the collapsibility relation $\sim$. Objects in the bottom level S_0 are collapsible when their attributes are identical. Objects at any level $\{S_l\}_{l=1}^L$ are collapsible if they are made of the same sets of parts for all the membership types π.

In Figure 1 we provide a pictorial representation of the domain compression of an $\mathcal{H}$-decomposition (EGNN, described in Section 5.2). On the left we show the $\mathcal{H}$-decomposition of a graph taken from the IMDB-BINARY dataset (see Section 5.1) together with its compressed version on the right.

In order to compress $\mathcal{H}$-decompositions we adapt the lifted linear programming technique proposed by (Mladenov et al., 2012) to the SAEN architecture. A matrix $M \in \mathbb{R}^{n \times p}$ with $m \leq n$ distinct rows can be decomposed as the product DM^{comp} where M^{comp} is a compressed version of M in which the distinct rows of M appear exactly once. The Boolean decompression matrix, D, encodes the collapsibility relation among the rows of M so that $D_{ij} = 1$ iff the i^{th} row of M falls in the equivalence class j of $\sim$. A pseudo-inverse C of D can be computed by dividing the rows of $D^\top$ by their sum (where $D^\top$ is the transpose of D).

Example 1 *If define matrix M as in Eq. 3 we notice that row 1 and 4 share the encoding $[0, 0, 0]$, rows 3 and 5 share the encoding $[1, 1, 0]$ while the encoding $[1, 0, 1]$ appears only once at row 2. So matrix M can be compressed to matrix M^{comp}.*

$$M = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 0 & 0 \\ 1 & 1 & 0 \end{bmatrix} \quad M^{comp} = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$
$$C = \begin{bmatrix} 1/2 & 0 & 0 & 1/2 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & 0 & 1/2 \end{bmatrix} \quad D = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (3)$$

Matrix M can be expressed as the matrix product between the decompression matrix D and the compressed version of M^{comp} (i.e. $M = DM^{comp}$), while the matrix multiplication between the compression matrix C and the M leads to the compressed matrix M^{comp} (i.e.$M^{comp} = CM$).

We apply domain compression to SAEN by rewriting Eqs. 1 and 2 in matrix form.

We rewrite Eq. 1 as:

$$A_l = \begin{cases} X & \text{if } l = 0 \\ \mathbf{R}_l \mathbf{H}_{l-1} & \text{if } l > 0 \end{cases} \quad (4)$$

where:

- $A_l \in \mathbb{R}^{|S_l| \times n(l-1)d(l)}$ is the matrix that represents the *shift-aggregated* vector representations of the object of level S_{l-1};
- $X \in \mathbb{R}^{|S_0| \times p}$ is the matrix that represents the p-dimensional encodings of the vertex attributes in V (i.e. the rows of X are the $\mathbf{x}_{v_i}$ of Eq. 1);

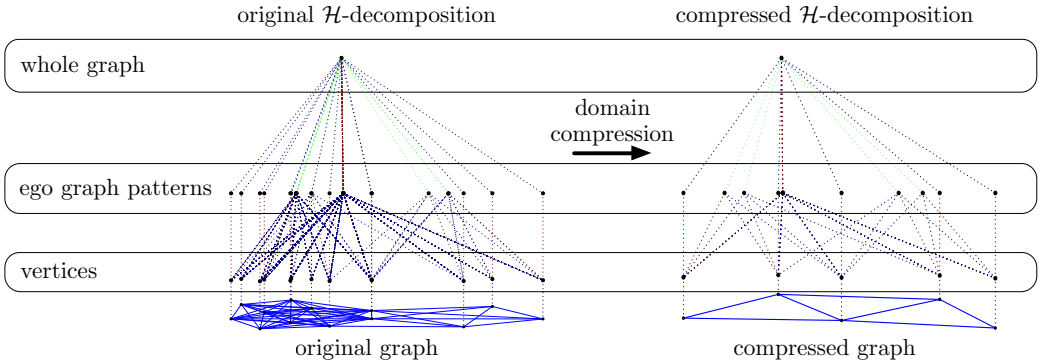

Figure 1. Pictorial representation of the $\mathcal{H}$-decomposition of a graph taken from the IMDB-BINARY dataset (see Section 5.1) together with its compressed version.

- $\mathbf{R}_l \in \mathbb{R}^{|S_l| \times n(l)|S_{l-1}|}$ is the concatenation

$$\mathbf{R}_l = \left[R_{l,1}, \dots, R_{l,\pi}, \dots, R_{l,n(l)}\right] \tag{5}$$

  of the matrices $R_{l,\pi} \in \mathbb{R}^{|S_l| \times |S_{l-1}|}\ \forall \pi \in \Pi_l$ which represent the $\mathcal{R}_{l,\pi}$-convolution relations of Eq. 1 whose elements are $(R_{l,\pi})_{ij} = 1$ if $(s', s) \in \mathcal{R}_{l,\pi}$ and 0 otherwise.

- $\mathbf{H}_{l-1} \in \mathbb{R}^{n(l)|S_{l-1}| \times n(l)d(l)}$ is a block-diagonal matrix

$$\mathbf{H}_{l-1} = \begin{bmatrix} H_{l-1} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & H_{l-1} \end{bmatrix} \tag{6}$$

  whose blocks are formed by matrix $H_{l-1} \in \mathbb{R}^{|S_{l-1}| \times d(l)}$ repeated $n(l)$ times. The rows of H_{l-1} are the vector representations $\mathbf{h}_j$ in Eq. 1.

Eq. 2 is simply rewritten to $H_l = f_l(A_l; \theta_l)$ where $f_l(\cdot; \theta_l)$ is unchanged w.r.t. Eq. 2 and is applied to its input matrix A_l row-wise.

Algorithm 1 DOMAIN-COMPRESSION

DOMAIN-COMPRESSION(X, R)
1 $C_0, D_0 =$ COMPUTE-CD(X)
2 $X^{comp} = C_0 X$
3 $R^{comp} = \{\}$
4 **for** $l = 1$ **to** L
5 $\quad R^{col_comp} = [R_{l,\pi} D_{l-1},\ \forall \pi = 1, \dots, n(l)]$
6 $\quad C_l, D_l =$ COMPUTE-CD(R^{col_comp})
7 $\quad$ **for** $\pi = 1$ **to** $n(l)$
8 $\quad\quad R^{comp}_{l,\pi} = C_l R^{col_comp}_{\pi}$
9 **return** X^{comp}, R^{comp}

Domain compression on Eq. 4 is performed by the DOMAIN-COMPRESSION procedure (see Algorithm 1). which takes as input the attribute matrix $X \in \mathbb{R}^{|S_0| \times p}$ and the part-of matrices $R_{l,\pi}$ and returns their compressed versions X^{comp} and the $R^{comp}_{l,\pi}$ respectively. The algorithm starts by invoking (line 1) the procedure COMPUTE-CD on X to obtain the compression and decompression matrices C_0 and D_0 respectively. The compression matrix C_0 is used to compress X (line 2) then we start iterating over the levels $l = 0, \dots, L$ of the $\mathcal{H}$-decomposition (line 4) and compress the $R_{l,\pi}$ matrices. The compression of the $R_{l,\pi}$ matrices is done by right-multiplying them by the decompression matrix D_{l-1} of the previous level $l - 1$ (line 5). In this way we collapse the parts of relation $\mathcal{R}_{l,\pi}$ (i.e. the columns of $R_{l,\pi}$) as these were identified in level S_{l-1} as identical objects (i.e. those objects corresponding to the rows of X or $R_{l-1,\pi}$ collapsed during the previous step). The result is a list $R^{col_comp} = [R_{l,\pi} D_{l-1},\ \forall \pi = 1, \dots, n(l)]$ of column compressed $R_{l,\pi}-$matrices. We proceed collapsing equivalent objects in level S_l, i.e. those made of identical sets of parts: we find symmetries in R^{col_comp} by invoking COMPUTE-CD (line 6) and obtain a new pair C_l, D_l of compression, and decompression matrices respectively. Finally the compression matrix C_l is applied to the column-compressed matrices in R^{col_comp} in order to obtain the Π_l compressed matrices of level S_l (line 8).

Algorithm 1 allows us to compute the domain compressed version of Eq. 4 which can be obtained by replacing: X with $X^{comp} = C_0 X$, $R_{l,\pi}$ with $R^{comp}_{l,\pi} = C_l R_{l,\pi} D_{l-1}$ and H_l with H^{comp}_l. Willing to recover the original encodings H_l we just need to employ the decompression matrix D_l on the compressed encodings H^{comp}_l, indeed $H_l = D_l H^{comp}_l$.

As we can see by substituting S_l with S^{comp}_l, the more are the symmetries (i.e. when $|S^{comp}_l| \ll |S_l|$) the greater the domain compression will be.

Table 1. Statistics of the datasets used in our experiments.

DATASET	SIZE	AVG. NODES	AVG. MAX. DEGREE
COLLAB	5000	74.49	73.62
IMDB-BINARY	1000	19.77	18.77
IMDB-MULTI	1500	13.00	12.00
REDDIT-BINARY	2000	429.62	217.35
REDDIT-MULTI5K	5000	508.51	204.08
REDDIT-MULTI12K	11929	391.40	161.70
MUTAG	188	17.93	3.01
PTC	344	25.56	3.73
NCI1	4110	29.87	3.34
PROTEINS	1113	39.06	5.79
D&D	1178	284.32	9.51

5. Experimental evaluation

We perform an exhaustive experimental evaluation and answer the following questions:

Q1 How does SAEN compare to the state of the art?

Q2 Can SAEN exploit symmetries in social networks to reduce the memory usage and the runtime?

5.1. Datasets

In order to answer the experimental questions we tested our method on six publicly available datasets first proposed by Yanardag & Vishwanathan (2015) and some bioinformatic datasets.

COLLAB is a dataset where each graph represent the ego-network of a researcher, and the task is to determine the field of study of the researcher between *High Energy Physics*, *Condensed Matter Physics* and *Astro Physics*.

IMDB-BINARY, IMDB-MULTI are datasets derived from IMDB where in each graph the vertices represent actors/actresses and the edges connect people which have performed in the same movie. Collaboration graphs are generated from movies belonging to genres *Action* and *Romance* for IMDB-BINARYand *Comedy*, *Romance* and *Sci-Fi* for IMDB-MULTI, and for each actor/actress in those genres an ego-graph is extracted. The task is to identify the genre from which the ego-graph has been generated.

REDDIT-BINARY, REDDIT-MULTI5K, REDDIT-MULTI12K are datasets where each graph is derived from a discussion thread from Reddit. In those datasets each vertex represent a distinct user and two users are connected by an edge if one of them has responded to a post of the other in that discussion. The task in REDDIT-BINARY is to discriminate between threads originating from a discussion-based subreddit (*TrollXChromosomes*, *atheism*) or from a question/answers-based subreddit (*IAmA*, *AskReddit*). The task in REDDIT-MULTI5K and REDDIT-MULTI12K is a multiclass classification problem where each graph is labeled with the subreddit where it has originated (*worldnews, videos, AdviceAnimals, aww, mildlyinteresting* for REDDIT-MULTI5K and *AskReddit, AdviceAnimals, atheism, aww, IAmA, mildlyinteresting, Showerthoughts, videos, todayilearned, worldnews, TrollXChromosomes* for REDDIT-MULTI12K).

MUTAG, PTC, NCI1, PROTEINS and D&D are bioinformatic datasets. MUTAG (Debnath et al., 1991) is a dataset of 188 mutagenic aromatic and heteroaromatic nitro compounds labeled according to whether or not they have a mutagenic effect on the Gramnegative bacterium *Salmonella typhimurium*. PTC (Toivonen et al., 2003) is a dataset of 344 chemical compounds that reports the carcinogenicity for male and female rats and it has 19 discrete labels. NCI1 (Wale et al., 2008) is a dataset of 4100 examples and is a subset of balanced datasets of chemical compounds screened for ability to suppress or inhibit the growth of a panel of human tumor cell lines, and has 37 discrete labels. PROTEINS (Borgwardt et al., 2005a) is a binary classification dataset made of 1113 proteins. Each protein is represented as a graph where nodes are secondary structure elements (i.e. helices, sheets and turns). Edges connect nodes if they are neighbors in the amino-acid sequence or in the 3D space. D&D is a binary classification dataset of 1178 graphs. Each graph represents a protein nodes are amino acids which are connected by an edge if they are less than 6 Angstroms apart.

5.2. Experiments

In our experiments we chose an $\mathcal{H}$-decomposition called Ego Graph Neural Network (EGNN) (shown in Figure 2), that mimics the graph kernel NSPDK with the distance parameter set to 0. Before applying EGNN we turn unattributed graphs (V, E) into attributed graphs (V, E, X) by annotating their vertices $v \in V$ with attributes $\mathbf{x}_v \in X$. We label vertices v of G with their degree and encode this information into the attributes $\mathbf{x}_v$ by employing the 1-hot encoding.

EGNN decomposes attributed graphs $G = (V, E, X)$ into a 3 level $\mathcal{H}$-decomposition with the following levels:

- level S_0 contains objects s_v that are in one-to-one correspondence with the vertices $v \in V$.

Table 2. Comparison of accuracy results on social network datasets.

DATASET	DGK (Yanardag & Vishwanathan, 2015)	PATCHY-SAN (Niepert et al., 2016)	SAEN (our method)
COLLAB	73.09 ± 0.25	72.60 ± 2.16	75.63 ± 0.31
IMDB-BINARY	66.96 ± 0.56	71.00 ± 2.29	71.26 ± 0.74
IMDB-MULTI	44.55 ± 0.52	45.23 ± 2.84	49.11 ± 0.64
REDDIT-BINARY	78.04 ± 0.39	86.30 ± 1.58	86.08 ± 0.53
REDDIT-MULTI5K	41.27 ± 0.18	49.10 ± 0.70	52.24 ± 0.38
REDDIT-MULTI12K	32.22 ± 0.10	41.32 ± 0.42	46.72 ± 0.23

Table 3. Parameters used for the EGNN decompositions for each datasets.

DATASET	RADIUSES	HIDDEN UNITS		
	r	S_0	S_1	S_2
COLLAB	$0, 1$	$15-5$	$5-2$	$5-3$
IMDB-BINARY	$0, 1, 2$	2	$5-2$	$5-3-1$
IMDB-MULTI	$0, 1, 2$	2	$5-2$	$5-3$
REDDIT-BINARY	$0, 1$	$10-5$	$5-2$	$5-3-1$
REDDIT-MULTI5K	$0, 1$	10	10	$6-5$
REDDIT-MULTI12K	$0, 1$	10	10	$20-11$
MUTAG	$0, 1, 2, 3$	10	$5-5$	$5-5-1$
PTC	$0, 1$	15	15	$15-1$
NCI1	$0, 1, 2, 3$	15	15	$15-10-1$
PROTEINS	$0, 1, 2, 3$	$3-2$	$6-5-4$	$6-3-1$
D&D	$0, 1, 2, 3$	10	$5-2$	$5-3-1$

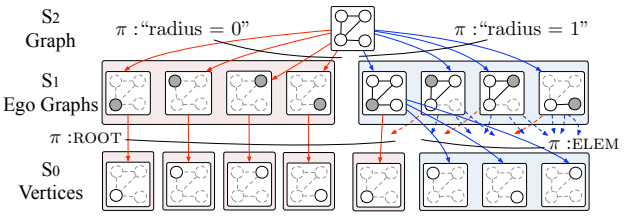

Figure 2. Example of Ego Graph decomposition.

- level S_1 contains v_{root}-rooted r-neighborhood subgraphs (i.e. ego graphs) $e = (v_{root}, V_e, E_e)$ of radius $r = 0, 1, \ldots, R$ and has part-of alphabet $\Pi_1 = \{\text{ROOT}, \text{ELEM}\}$. Objects $s_v \in S_0$ are "ELEM-part-of" ego graph e if $v \in V_e \setminus \{v_{root}\}$, while the are "ROOT-part-of" ego graph e if $v = v_{root}$.
- level S_2 contains the graph G that we want to classify and has part-of alphabet $\Pi_2 = \{0, 1\}$ which correspond to the radius of the ego graphs $e \in S_1$ of which G is made of.

The EGNN decomposition is exemplified for a small graph shown in Figure 2.

E1 We experimented with SAEN applying the EGNN $\mathcal{H}$-decomposition on all the datasets. For each dataset, we manually chose the parameters of SAEN, i.e. the number of hidden layers for each level, the size of each layer and the maximum radius R. We used the Leaky ReLU (Maas et al., 2013) activation function on all the units. In Table 3 we report for each dataset the radiuses r of the neighborhood subgraphs used in the EGNN decomposition and the number of units in the hidden layers for each level.

In all our experiments we trained the neural networks by using the Adam algorithm to minimize a cross entropy loss.

The classification accuracy of SAEN was measured with 10-times 10-fold cross-validation. We manually chose the number of layers and units for each level of the part-of decomposition; the number of epochs was chosen manually for each dataset and we kept the same value for all the 100 runs of the 10-times 10-fold cross-validation.

The mean accuracies and their standard deviations obtained by applying our method on social network datasets are reported in Table 2, where we compare these results with those obtained by Yanardag & Vishwanathan (2015) and by Niepert et al. (2016). In Table 5 we compare the results obtained by our method on bioinformatic datasets with the those obtained by Niepert et al. (2016).

Table 4. Comparison of sizes and runtimes of the datasets before and after the compression.

DATASET	SIZE (MB)			RUNTIME		
	ORIGINAL	COMP.	RATIO	ORIGINAL	COMP.	SPEEDUP
COLLAB	1190	448	0.38	43' 18"	8' 20"	5.2
IMDB-BINARY	68	34	0.50	3' 9"	0' 30"	6.3
IMDB-MULTI	74	40	0.54	7' 41"	1' 54"	4.0
REDDIT-BINARY	326	56	0.17	TO	2' 35"	≥ 100.0
REDDIT-MULTI5K	952	162	0.17	OOM	9' 51"	–
REDDIT-MULTI12K	1788	347	0.19	OOM	29' 55"	–

Table 5. Comparison of accuracy on bio-informatics datasets.

DATASET	PATCHY-SAN (Niepert et al., 2016)	SAEN (our method)
MUTAG	92.63 ± 4.21	84.99 ± 1.82
PTC	62.29 ± 5.68	57.04 ± 1.30
NCI1	78.59 ± 1.89	77.80 ± 0.42
PROTEINS	75.89 ± 2.76	75.31 ± 0.70
D&D	77.12 ± 2.41	77.69 ± 0.96

E2 In Table 4 we show the file sizes of the preprocessed datasets before and after the compression together with the data compression ratio. [2] We also estimate the benefit of the relational compression from a computational time point of view and report the measurement of the runtime for 1 run with and without compression together with the speedup factor.

For the purpose of this experiment, all tests were run on a computer with two 8-cores Intel Xeon E5-2665 processors and 94 GB RAM. Uncompressed datasets which exhausted our server's memory during the test are marked as "OOM" (out of memory) in the table, while those who exceeded the time limit of 100 times the time needed for the uncompressed version are marked as "TO" (timeout). SAEN was implemented in Python with the TensorFlow library.

5.3. Discussion

A1 As shown in Table 2, EGNN performs consistently better than the other two methods on all the social network datasets. This confirms that the chosen $\mathcal{H}$-decomposition is effective on this kind of problems. Table 1 shows that the average maximum node degree (AMND) [3] of the social network datasets is in the order of 10^2. SAEN can easily cope with highly skewed node degree distributions by aggregating distributed representation of patterns while this is not the case for DGK and PATCHY-SAN. DGK uses the same patterns of the corresponding non-deep graph kernel used to match common substructures. If the pattern distribution is affected by the degree distribution most of those patterns will not match, making it unlikely for DGK to work well on social network data. PATCHY-SAN employs as patterns neighborhood subgraphs truncated or padded to a size k in order to fit the size of the receptive field of a CNN. However, since Niepert et al. (2016) experiment with $k = 10$, it is not surprising that they perform worst than SAEN on COLLAB, IMDB-MULTI, REDDIT-MULTI5K and REDDIT-MULTI12K since a small k causes the algorithm to throw away most of the subgraph; a more sensible choice for k would have been the AMND of each graph (i.e. 74, 12, 204 and 162 respectively, cf. Tables 1 and 2).

Table 5 compares the results of SAEN with the best PATCHY-SAN instance on chemoinformatics and bioinformatics datasets. SAEN is in line with the results of Niepert et al. (2016) on PROTEINS and D&D, two datasets where the degree is in the order of 10 (see Table 1). Small molecules, on the other hand, have very small degrees. Indeed, in NCI1, MUTAG and PTC SAEN does not perform very well and is outperformed by PATCHY-SAN, confirming that SAEN is best suited for graphs with large degrees. Incidentally, we note that for small molecules, graph kernels attain even better accuracies (e.g. the Weisfeiler-Lehman graph kernel (Shervashidze et al., 2011) achieves 80.13% accuracy on NCI1).

A2 The compression algorithm has proven to be effective in improving the computational cost of our method. Most of the datasets improved their runtimes by a factor of at least 4 while maintaining the same expressive power. Moreover, experiments on REDDIT-MULTI5K and REDDIT-MULTI12K have only been possible thanks to the size reduction operated by the algorithm as the script exhausted the memory while executing the training step on the uncompressed files.

[2] The size of the uncompressed files are shown for the sole purpose of computing the data compression ratio. Indeed the last version of our code compresses the files on the fly.

[3] The AMND for a given dataset is obtained by computing the maximum node degree of each graph and then averaging over all graphs.

6. Related works

When learning with graph inputs two fundamental design aspects that must be taken into account are: the choice of the pattern generator and the choice of the matching operator. The former decomposes the graph input in substructures while the latter allows to compare the substructures.

Among the patterns considered from the graph kernel literature we have paths, shortest paths, walks (Kashima et al., 2003), subtrees (Ramon & Gärtner, 2003; Shervashidze et al., 2011) and neighborhood subgraphs (Costa & De Grave, 2010). The similarity between graphs G and G' is computed by counting the number of matches between their common the substructures (i.e. a kernel on the sets of the substructures). The match between two substructures can be defined by using graph isomorphism or some other weaker graph invariant.

When the number of substructures to enumerate is infinite or exponential with the size of the graph (perhaps this is the case for random walks and shortest paths respectively) the kernel between the two graphs is computed without generating an explicit feature map. Learning with an implicit feature map is not scalable as it has a space complexity quadratic in the number of training examples (because we need to store in memory the gram matrix).

Other graph kernels such as the Weisfeiler-Lehman subtree kernel (WLST) (Shervashidze et al., 2011) and the Neighborhood Subgraph Pairwise Distance Kernel (NSPDK) (Costa & De Grave, 2010) deliberately choose a pattern generator that scales polynomially and produces an explicit feature map. However the vector representations produced by WLST and NSPDK are handcrafted and not learned.

Deep graph kernels (DGK) (Yanardag & Vishwanathan, 2015) upgrade existing graph kernels with a feature reweighing schema. DGKs represent input graphs as a corpus of substructures (e.g. graphlets, Weisfeiler-Lehman subtrees, vertex pairs with shortest path distance) and then train vector embeddings of substructures with CBOW/Skip-gram models. [4] Each graph-kernel feature (i.e. the number of occurrences of a substructure) is reweighed by the 2-norm of the vector embedding of the corresponding substructure. Experimental evidence shows that DGKs alleviate the problem of diagonal dominance in GKs. However, DGKs inherit from GKs a flat representation (i.e. just one layer of depth) of the input graphs and the vector representations of the substructures are not trained end-to-end as SAEN would do.

PATCHY-SAN (Niepert et al., 2016) casts graphs into a format suitable for learning convolutional neural networks (CNNs): 1) graphs are decomposed into a fixed number of neighborhood subgraphs; 2) which are then casted to a fixed-size receptive field. Both 1) and 2) involve either padding or truncation in order to meet the fixed-size requirements. The truncation operation can be detrimental for the statistical performance of the downstream CNN since it throws away part of the input graph. On the other hand SAEN is able to handle structured inputs of variable sizes without throwing away part of the them. And this is one of the reasons because SAEN has better statistical performance than PATCHY-SAN (See § 5).

[4] The CBOW/Skip-gram models receive as inputs cooccurrences among substructures sampled from the input graphs.

7. Conclusions

Hierarchical decompositions introduce a novel notion of depth in the context of learning with structured data, leveraging the nested part-of-parts relation. In this paper, we defined a simple architecture based on neural networks for learning representations of these hierarchies. We showed experimentally that the approach is particularly well-suited for dealing with graphs that are large and have high degree, such as those that naturally occur in social network data. Our approach is also effective for learning with smaller graphs, such as those occurring in chemoinformatics and bioinformatics, although in these cases the performance of SAEN does not exceed the state-of-the-art established by other methods. A second contribution of this paper is the domain compression algorithm, which greatly reduces memory usage and allowed us to speedup the training time of a factor of at least 4.

References

Baldi, P and Pollastri, G. The principled design of large-scale recursive neural network architectures–dag-rnns and the protein structure prediction problem. *J Mach Learn Res*, 4(Sep):575–602, 2003.

Bianucci, Anna Maria, Micheli, Alessio, Sperduti, Alessandro, and Starita, Antonina. Application of cascade correlation networks for structures to chemistry. *Applied Intelligence*, 12(1-2):117–147, 2000.

Borgwardt, K M and Kriegel, H-P. Shortest-path kernels on graphs. In *Proc. of the ICDM-05*, pp. 8–pp. IEEE, 2005.

Borgwardt, Karsten M, Ong, Cheng Soon, Schönauer, Stefan, Vishwanathan, SVN, Smola, Alex J, and Kriegel, Hans-Peter. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl 1):i47–i56, 2005a.

Borgwardt, Karsten M., Ong, Cheng Soon, Schnauer, Stefan, Vishwanathan, S. V. N., Smola, Alex J., and Kriegel, Hans-Peter. Protein function prediction via

graph kernels. *Bioinformatics*, 21(suppl 1):i47–i56, 2005b. 00313.

Costa, F and De Grave, K. Fast neighborhood subgraph pairwise distance kernel. In *Proc. of the ICML-10*, pp. 255–262. Omnipress, 2010.

Debnath, Asim Kumar, Lopez, de Compadre RL, Debnath, Gargi, Shusterman, Alan J, and Hansch, Corwin. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *J. Med. Chem.*, 34(2):786–797, 1991.

Goller, C and Kuchler, A. Learning task-dependent distributed representations by backpropagation through structure. In *Neural Networks, 1996., IEEE International Conference on*, volume 1, pp. 347–352. IEEE, 1996.

Haussler, D. Convolution kernels on discrete structures. Technical report, Citeseer, 1999.

Kashima, H, Tsuda, K, and Inokuchi, A. Marginalized kernels between labeled graphs. In *n Proc. of the ICML-03*, volume 3, pp. 321–328, 2003.

Maas, Andrew L, Hannun, Awni Y, and Ng, Andrew Y. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, volume 30, 2013.

Mladenov, M, Ahmadi, B, and Kersting, K. Lifted linear programming. In *AISTATS-12*, pp. 788–797, 2012.

Niepert, M, Ahmed, M, and Kutzkov, K. Learning convolutional neural networks for graphs. *arXiv preprint arXiv:1605.05273*, 2016.

Ralaivola, Liva, Swamidass, Sanjay J., Saigo, Hiroto, and Baldi, Pierre. Graph kernels for chemical informatics. *Neural Networks*, 18(8):1093–1110, 2005. 00251.

Ramon, J and Gärtner, T. Expressivity versus efficiency of graph kernels. In *First International Workshop on Mining Graphs, Trees and Sequences*, pp. 65–74. Citeseer, 2003.

Scarselli, Franco, Gori, Marco, Tsoi, Ah Chung, Hagenbuchner, Markus, and Monfardini, Gabriele. The graph neural network model. *Neural Networks, IEEE Transactions on*, 20(1):61–80, 2009. 00073.

Schoelkopf, Bernhard, Weston, Jason, Eskin, Eleazar, Leslie, Christina, and Noble, William Stafford. A kernel approach for learning from almost orthogonal patterns. In *European Conference on Machine Learning*, pp. 511–528. Springer, 2002.

Shervashidze, N, Vishwanathan, SVN, Petri, T, Mehlhorn, K, and Borgwardt, K M. Efficient graphlet kernels for large graph comparison. In *AISTATS-09*, volume 5, pp. 488–495, 2009.

Shervashidze, N, Schweitzer, P, Leeuwen, E J van, Mehlhorn, K, and Borgwardt, K M. Weisfeiler-lehman graph kernels. *J Mach Learn Res*, 12(Sep):2539–2561, 2011.

Socher, R, Lin, C C, Manning, C, and Ng, A Y. Parsing natural scenes and natural language with recursive neural networks. In *Proc. of the ICML-11*, pp. 129–136, 2011.

Sperduti, A and Starita, A. Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(3):714–735, 1997.

Toivonen, Hannu, Srinivasan, Ashwin, King, Ross D, Kramer, Stefan, and Helma, Christoph. Statistical evaluation of the predictive toxicology challenge 2000–2001. *Bioinformatics*, 19(10):1183–1193, 2003.

Vullo, A and Frasconi, P. Disulfide connectivity prediction using recursive neural networks and evolutionary information. *Bioinformatics*, 20(5):653–659, 2004.

Wale, Nikil, Watson, Ian A, and Karypis, George. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3):347–375, 2008.

Yanardag, P and Vishwanathan, SVN. Deep graph kernels. In *Proc. of KDD-15*, pp. 1365–1374, 2015.